

CLAIMS

- 1 1. Apparatus for performing speculative prefetching for a PCI DMA read request in
2 a PCI - InfiniBand bridge system, the apparatus comprising:
 - 3 an update mechanism responsive to data returning from an initial
4 InfiniBand RDMA request issued to satisfy the PCI DMA read request for
5 computing a next address from the read address of the initial InfiniBand RDMA
6 request; and
 - 7 a repeat mechanism that generates a new RDMA read request using the
8 next address to perform a speculative prefetch.
- 1 2. The apparatus of claim 1 wherein the next address is computed from information
2 defining a prefetch request.
- 1 3. The apparatus of claim 2 wherein the information defining a prefetch request
2 comprises a prefetch request number and a prefetch request size.
- 1 4. The apparatus of claim 3 wherein the update mechanism computes the next
2 address by adding the prefetch request size to the read address of the initial
3 InfiniBand RDMA request.
- 1 5. The apparatus of claim 3 wherein the repeat mechanism generates further
2 RDMA prefetch read requests until an amount of data equal to the prefetch
3 request size has been retrieved.
- 1 6. The apparatus of claim 1 further comprising a data structure that stores a data
2 tag identifying the new RDMA read request.

- 1 7. The apparatus of claim 6 wherein the data structure is referenced in the new
- 2 RDMA read request.

- 1 8. The apparatus of claim 1 wherein the repeat mechanism generates the new
- 2 RDMA read request on a work queue used to generate the initial RDMA read
- 3 request.

- 1 9. The apparatus of claim 1 wherein the data returning from an RDMA read request
- 2 comprises a plurality of data packets.

- 1 10. The apparatus of claim 1 wherein the update mechanism comprises an address
- 2 map that computes the read address of the initial InfiniBand RDMA request from
- 3 a PCI address.

- 1 11. The apparatus of claim 10 wherein the address map includes an R-key that is associated with an area in a memory from which DMA data is retrieved and a pointer to a work queue that generates the initial InfiniBand RDMA request.

- 1 12. The apparatus of claim 11 wherein the repeat mechanism generates a new read address for an RDMA read request by combining the PCI address with the R-Key and the next address.

- 1 13. A method for performing speculative prefetching for a PCI DMA read request in a PCI - InfiniBand bridge system, the method comprising:
 - 3 (a) in response to data returning from an initial InfiniBand RDMA request issued to satisfy the PCI DMA read request, computing a next address from the read address of the initial InfiniBand RDMA request; and
 - 6 (b) generating a new RDMA read request using the next address to perform a speculative prefetch.

- 1 14. The method of claim 13 wherein step (b) comprises computing the next address
 - 2 from information defining a prefetch request.
 - 1 15. The method of claim 14 wherein the information defining a prefetch request
 - 2 comprises a prefetch request number and a prefetch request size.
 - 1 16. The method of claim 15 wherein step (a) comprises computing the next address
 - 2 by adding the prefetch request size to the read address of the initial InfiniBand
 - 3 RDMA request.
 - 1 17. The method of claim 15 wherein step (b) comprises generating further RDMA
 - 2 prefetch read requests until an amount of data equal to the prefetch request size
 - 3 has been retrieved.
 - 1 18. The method of claim 13 further comprising (c) storing a data tag identifying the
 - 2 new RDMA read request in a data structure.
 - 1 19. The method of claim 18 wherein the data structure is referenced in the new
 - 2 RDMA read request.
 - 1 20. The method of claim 13 wherein step (b) comprises generating the new RDMA
 - 2 read request on a work queue used to generate the initial RDMA read request.
 - 1 21. The method of claim 13 wherein the data returning from an RDMA read request
 - 2 comprises a plurality of data packets.

- 1 22. The method of claim 13 wherein the step (a) comprises using an address map to
- 2 compute the read address of the initial InfiniBand RDMA request from a PCI
- 3 address.
- 1 23. The method of claim 22 wherein the address map includes an R-key that is associated with an area in a memory from which DMA data is retrieved and a pointer to a work queue that generates the initial InfiniBand RDMA request.
- 1 24. The method of claim 23 wherein step (b) comprises generating a new read address for an RDMA read request by combining the PCI address with the R-Key and the next address.
- 1 25. Apparatus for performing speculative prefetching for a PCI DMA read request in a PCI - InfiniBand bridge system, the apparatus comprising:
 - 3 means responsive to data returning from an initial InfiniBand RDMA request issued to satisfy the PCI DMA read request for computing a next address from the read address of the initial InfiniBand RDMA request; and
 - 6 means for generating a new RDMA read request using the next address to perform a speculative prefetch.
- 1 26. The apparatus of claim 25 wherein the means for computing the next address comprises means for computing the next address from information defining a prefetch request.
- 1 27. The apparatus of claim 26 wherein the information defining a prefetch request comprises a prefetch request number and a prefetch request size.

- 1 33. The method of claim 32 wherein step (a) comprises incrementing the current
 - 2 address by adding the prefetch request size to the current address.
 - 1 34. The method of claim 33 wherein step (a) comprises using the DMA scoreboard to
 - 2 store a data tag identifying the prefetch read request.
 - 1 35. The method of claim 31 wherein step (b) comprises generating a new DMA read
 - 2 request on a work queue used to generate the initial DMA read request.
 - 1 36. The method of claim 35 wherein step (a) comprises referencing the DMA
 - 2 scoreboard data structure in the new DMA read request.
 - 1 37. The method of claim 31 wherein step (b) comprises generating further DMA
 - 2 prefetch read requests until an amount of data equal to the prefetch size has
 - 3 been retrieved.
 - 1 38. The method of claim 31 wherein the response to a DMA read request comprises
 - 2 a plurality of data packets.
 - 1 39. The method of claim 38 wherein the current address is a PCI address.
 - 1 40. The method of claim 39 wherein step (b) comprises generating a new DMA read
 - 2 request by combining the PCI address with an R-Key that is associated with an
 - 3 area in the memory.